



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ :

G06F 17/30

A2

(11) International Publication Number:

WO 97/40456

(43) International Publication Date:

30 October 1997 (30.10.97)

(21) International Application Number: PCT/US97/06183

(22) International Filing Date: 24 April 1997 (24.04.97)

(30) Priority Data:

08/639,073	24 April 1996 (24.04.96)	US
08/825,276	27 March 1997 (27.03.97)	US

(71) Applicant: LINGUATEQ, INC. [US/US]; 1359 Beverly Road, McClean, VA 22101 (US).

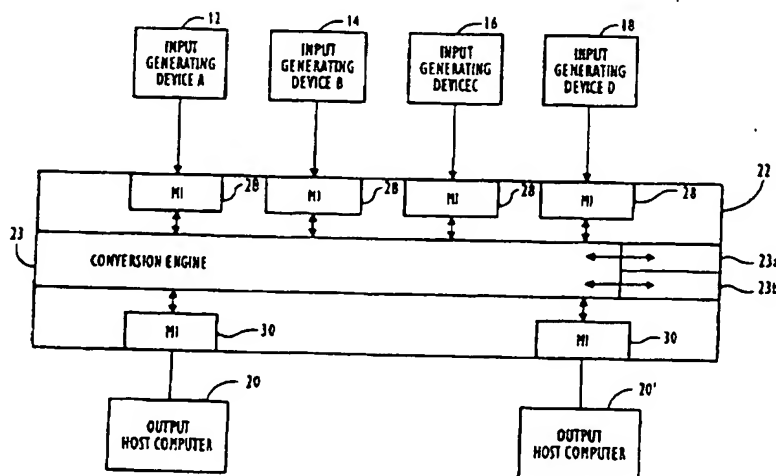
(72) Inventors: WILLIAMS, Dale, L.; 321 Springvale Road, Great Falls, VA 22066 (US). PETERSON, Michael, E.; 6821 Broyhill Street, McLean, VA 22101 (US).

(74) Agents: DEVINSKY, Paul et al.; McDermott, Will & Emery, 1850 K Street, N.W., Washington, DC 20006 (US).

(81) Designated States: AM, AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LT, LU, LV, MD, MG, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TT, UA, UG, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

Published*Without international search report and to be republished upon receipt of that report.*

(54) Title: METHOD AND APPARATUS FOR RATIONALIZING DIFFERENT DATA FORMATS IN A DATA MANAGEMENT SYSTEM



(57) Abstract

A record conversion device and method, including a preferably software implemented conversion engine, which receives formatted input records from an input record generating device, converts the records into a different (universal) format in real time and, depending upon the output host device(s), into an output specific format and transmits the converted records to the output host device(s). Messages, i.e., information, from output host device(s) can also be transferred, via the conversion device, to the input record generating devices. By rapidly processing new and/or incompatible record formats, the invention minimizes the cost of maintaining information management systems and maximizes their usage. The conversion device input and output architecture is flexible and the conversion device can be used with multiple (networked or non-networked) input record generating devices and single or multiple record storage devices using multiple record output formats.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon		Republic of Korea	PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

Method and Apparatus For
Rationalizing Different Data Formats
In a Data Management System

Field of the Invention

The present invention is a continuation-in-part of Application Serial Number 08/639,073 filed April 24, 1996 and relates to a method and apparatus for rationalizing and/or routing data records sent from various generating devices and/or devices using diverse data record formats, protocols and the like, and destined for further devices and further processing. More specifically, the present invention relates to a method and apparatus for eliminating the need to modify a data record processing system each time a new record input device, record file, or data record processing method is introduced. Still more specifically, the present invention relates to a method and device for generating desired data record outputs in a desired format irrespective of the format or protocols of the data record input device used.

Background

The complexity and effort required to manipulate the variety of unique records generated by present record generating devices has caused information system developers to develop unique (in the sense of customized) information systems and processes for each data record source. Current information gathering and processing systems, for example telecommunications information support systems, frequently receive data records from several input devices, many of which may use different record formats. In order to be able to manipulate the data in the records from these diverse sources, this information

processing environment requires that a custom interface exist² between each data record generating device and the data processing system. Fig. 1 depicts a typical current operating environment for an exemplary telecom billing system.

In the example illustrated in Fig. 1 each input device or data generator 12, 14, 16, 18 requires a custom interface 12', 14', 16', 18' respectively to input the data records it generates to a host computer 20 for processing. Similarly the host computer requires a custom interface 12', 14', 16' and 18' respectively for accepting data from each of the data generating devices. These are not trivial interfaces to design and implement and especially on the host computer side of the system, may represent a significant resource allocation. In the telecom example, assume that a telecom customer initially subscribes only to telephone service. Each transaction event, i.e. telephone call, triggers a data record reflecting that transaction in a specified format to be processed. If, over time, that customer adds cable TV to its service, a new and usually differently formatted stream of data records is generated each time a transaction event occurs. If still later, that customer adds a pager and a fax machine to its service, additional new and again differently formatted data record streams are generated for each transaction originating from those sources. All of these situations create potential problems of communication and data/record interoperability between the record input device and the information processing system.

It is also known from U.S. Patent No. 5,497,491 to provide object i.e. data structure, oriented computing environments in

which data is imported or exported³ using a common external interface for data structures within a data structure oriented environment by encapsulating the data using a "tagging" scheme in order to obviate the need for separate interfaces in order to communicate between the object oriented computing environment and other, external computing environment.. This technique does not generally involve the use logical conversion or validation operations.

The invention of U.S. Patent No. 5,497,491 is directed to a method of simplifying coding, i.e. software for mapping pairs of objects defined at program generation. As such, it is directed to the data processor side of a prior art system such as that described above.

Vendor generated problems, i.e., communications (or network) interoperability, are usually addressed by means of a network protocol convertor or by open standards. However, (again reverting to the telecom example) each time a new device or service is added (e.g., caller ID, call waiting, voice mail), a new, usually custom designed data/record format interface must also be designed in order to support that device or service. As more new telecom services, such as multiway pagers, home (electronic) banking, personal (Internet) home pages, and E-mail proliferate, and as industries such as the telecom industry experience consolidation and deregulation events where different corporate entities using different transaction data record formats or protocols merge and/or acquire each other, the problem of rationalizing record collection and data processing will, based on current technology, become more complex and severe.

4

The standard conventional solution, that of modifying or recreating an entire accounting or billing system to support that new service or device or creating a patchwork of customized data record formats become less tenable and less satisfactory as options. Also, for specialized data record streams (such as corporate transactions) industry wide standards are not a viable option as data record structures are often a vital part of a company's proprietary data processing operations.

Objects and Summary of the Invention

It is an object of the present invention to provide a method and apparatus for converting data records from one format into a desired format regardless of the format in which it was generated.

It is a further object of the invention to provide a method and apparatus for converting one or more input record formats into a common, universal format.

It is yet a further object of the invention to provide a method and apparatus for converting records from a universal format to one or more predetermined output formats.

It is yet a further object of the invention to provide a method and apparatus compatible with existing records processing systems to "translate" between multiple record formats.

It is still a further object of the invention to provide a method and apparatus which can be readily integrated into existing data processing systems and which can instantly adopt such systems to process and output appropriately formatted records as new products and/or services are introduced.

It is yet a further object⁵ of the invention to provide a method and apparatus for executing communication protocols between a record generator and host computer which processes the records, and in which the apparatus may be deployed centrally (i.e., at the host computer) or remotely (i.e., at the record generator).

Additional objects, advantages and novel features of the invention will be set forth in part in the description which follows, and in part will become apparent to those skilled in the art upon examination of the following or may be learned by practice of the invention. The objects and advantages of the invention may be realized and attained by means of the instrumentalities and combinations particularly pointed out in the appended claims.

To achieve the foregoing and other objects and in accordance with the purpose of the invention, as embodied and broadly described herein, the method and apparatus of this invention may comprise an information retrieval and transit storage system which comprises: input record generating device(s) for generating data records of a predetermined format(s); a transaction engine for receiving records from the input record generating device(s), converting the data records into a universal format, and from the universal format into a predetermined output format(s); and an output host device(s) for receiving records in the predetermined output format(s). The system preferably includes a plurality of input record generating devices, each connected to the transaction engine; the plurality of input record generating devices may be either networked or stand alone. The transaction engine preferably

comprises at least two input interfaces⁶, each supporting a different predetermined input record format and each being operable to convert the input records into a universal format. The transaction engine will preferably further comprises at least one output interface operable to convert records from the universal format into a predetermined output format(s).

In another aspect of the invention, there is provided a transaction engine for receiving data records in different formats from a plurality of input record generating devices. The transaction comprises at least one input interface for receiving incoming data records, the input interface supporting a predetermined input format; a conversion device for receiving input records from the input interface and converting the records into a predetermined universal format; and at least one output interface for receiving a record(s) in the universal format and forwarding it in a predetermined output format(s). The conversion device preferably operates in real-time and preferably includes a first storage device for temporary storage of records during conversion and a library of conversion rules for converting data records from one or more input formats into the universal format and from the universal format into one or more output formats.

In yet another aspect, the invention comprises a method of processing records from one or more input formats into one or more output formats using a conversion device, the method comprises the steps of: receiving records organized in a first format; converting the records in a first interface device into a universal format; converting the records in a second interface device from the universal format into a desired output

format; and transmitting the records in the output format to an output host device, e.g., a data processing computer.

Brief Description of the Drawings

The accompanying drawings, which are incorporated in and form a part of the specification, illustrate the presently preferred embodiments of the present invention, and, together with the description, serve to explain the principles of the invention. In the drawings:

Fig. 1 is a depiction of a prior art data management system utilizing custom designed interface devices for each type of input record generating device used.

Fig. 2A is a depiction of a record management system incorporating a transaction engine according to the present invention for each input record generating device.

Fig. 2B is a depiction of a management system incorporating a common transaction engine according to the present invention for a plurality of record host devices.

Fig. 2C is a depiction of a record management system incorporating a transaction engine connected to a plurality of data generating devices via a network and to a plurality of record host devices via a network.

Fig. 3A is a schematic depiction of the architecture of an exemplary transaction engine in simplified form for ease of explanation and illustration.

Fig. 3B is a depiction of the relationship between the input and output interface of the transaction engine of Fig. 3A, for explaining the relationship between the input and output

interfaces and the library and ⁸super-record of the transaction engine.

Fig. 4 is a schematic depiction of an exemplary management system including an input record generating device connected to a transaction engine.

Fig. 5 is a depiction of an exemplary management system according to the present invention in which a plurality of input record generating devices are connected to one transaction engines using a shared port connection.

Fig. 6 is a depiction of yet another exemplary embodiment of a record management system according to the present invention in which an input record generating device is connected to a transaction engine from which records are distributed to multiple output host devices.

Fig. 7 is a depiction of yet another exemplary embodiment of the invention in which two or more input record generating devices are connected to a transaction engine and from which records from each input record generator may be distributed to one or to multiple output host devices.

Fig. 8 is a simplified version of Fig. 3B illustrating information flow between an exemplary input record generating device, a transaction engine according to the present invention and an output host device, typically a computer for maintaining a data base.

Fig. 9A is a flow diagram illustrating the transaction processor control over the flow of messages and records from a generating device to the transaction engine of Fig. 8 during a transaction for exemplary types of transactions.

Fig. 9B is a flow diagram⁹ illustrating the flow of response messages from the transaction engine to an input record generating device in connection with a response message from an output host device(s) provoked by receipt of a converted transaction record.

Fig. 10A is a flow diagram illustrating the flow of records/messages between a specific input side transaction processor and output side transaction processor of the transaction engine of Fig. 9A during a transaction.

Fig. 10B is a flow diagram illustrating the flow of response data from a specific output side transaction processor to a specific input side transaction processor of the transaction engine of Fig. 9B in connection with a response from a downstream host device to receipt of a converted transaction record.

Fig. 11A is a flow diagram illustrating the preferred multiprocessing in the transaction engine on a decomposed task basis of a record received from a record generating device.

Fig. 11B is a flow diagram of the flow of a converted transaction record to a port associated with a specific output host device.

Detailed Description of the Invention

Reference will now be made in detail to the presently preferred embodiments of the invention, examples of which are illustrated in the accompanying drawings. In the various drawings, the same reference numerals have been used to refer to corresponding devices or elements.

Turning first to Fig. 1, in a typical prior art business system, such as that illustrated in Fig. 1, records may be received from several different input devices, for example a cellular telephone, 12, a cable TV, 14, a modem, 16, and facsimile machine, 18, each of which generates data records in a different format. As will be apparent to the artisan, the generating devices may also be various point of sale terminals or virtually any type of device which generates a stream of data records for further processing. In such an environment, specially designed, custom interface devices 12', 14', 16' and 18' must be created for communication between the respective generating devices, 12, 14, 16 and 18 and a host device 20, typically a data processing host computer where data records are processed. Similarly the host device is provided with interfaces 12', 14', 16' and 18' respectively for receiving record streams from each of the input devices. The need for custom interface devices not only increases cost and decreases efficiency, it results in the need to constantly redesign record formats, and design new custom interfaces, i.e., as new services are offered by the business involved or new record formats are introduced. Developing record formats and custom interfaces frequently requires allocation of significant resources and may also result in significant delays as new services and/or formats are introduced.

Fig. 2A depicts an exemplary system according to the present invention in which, like the system of Fig. 1, data in different record formats is generated by several different generating devices, 12, 14, 16 and 18. However, the system of Fig. 2A uses a conversion device in the form of a transaction

engine 22 according to the present invention which eliminates the need to modify the system each time a new record generating device, service or the like which uses a new or different data record format is introduced. In the system of Fig. 2A, the transaction engine 22 receives formatted data records from an input record generating device and ultimately converts those records into a desired format which may then be transmitted or otherwise distributed to the output host device(s) 20 and/or 20' for processing. Businesses such as those specializing in global telecommunications, credit card processing, electric utility service, and cable television can advantageously use a transaction engine 22 based system to quickly accommodate new business services. By rapidly, i.e. essentially instantaneously, processing new or incompatible transaction record formats, the transaction engine 22 minimizes the cost of maintaining existing business systems while maximizing the effective usage of those systems.

Fig. 2A is an example of a system in which the transaction engine device 22 is deployed remotely, i.e., at the record generator, or in proximity to two or more record generators. The transaction engine 22 may be connected directly to the output host device or may be connected by a network.

Alternatively, as illustrated in Fig. 2B, the input transaction engine 22 may be centrally deployed, i.e., at or in the vicinity of the output host device 20. In the case of Fig. 2B, the record generating devices are connected directly to the transaction engine 22. Alternatively, as illustrated in Fig. 2C, the input record generating devices may be so connected via

a network N1 and the host devices ¹²20 and 20' may be connected to the transaction engine 22 via a network, N2.

As best seen in Fig. 3A, the transaction engine 22 uses a unique message interface (MI) to process each supported record format. In a preferred implementation, when the transaction engine 22 receives a record, it will, as described in more detail below, control the appropriate MI by means of the conversion engine 23, which is preferably embedded in the transaction engine 22, to "translate" the record from the input format to a desired Universal Data Format (UDF), which may be an industry specific format. As described in more detail below with reference to Fig. 3B, the conversion engine 23, preferably includes a super-record storage 23b for holding records during the conversion processes and, a library storage 23a for storage of the record conversion rules applicable to the various message interface devices. More specifically, it is the library 23a which contains the conversion rules applicable to the field of the UDF.

When a conversion from an input format to UDF is complete, the conversion engine 23 passes control to an output MI 30 that again uses the conversion engine 23 to translate the record from the UDF to a specified output record format. As explained in more detail below, the conversion engine 23 uses separate conversion rules to translate a record from the input format to the UDF and from the UDF to the output format.

As will be understood by one of skill in the art, the conversion engine 23, which is preferably software driven, may operate using any appropriate operating system although the UNIX™ operating system and Sun® platform, such as for example a

platform using a Sun[®] chip set, ¹³ UltraSPARC™, are presently preferred.

As alluded to above, the input generating devices may be stand alone devices, nodes of a networked system or some combination of both. Although physical network interface specifications may vary from application to application, it is presently preferred that each transaction engine 22 should preferably support a plurality of interfaces such as the V.3X - type, 10BaseT, FDDI, or RJ-11/45 interfaces. The conversion engine 23 should preferably also support all major communication protocols such as for example the TCP/IP protocol suite, or the IBM[®] SNA™ protocol suite.

Referring now to Fig. 3B, in order to make the transaction engine 22 as flexible as possible, i.e., so as to meet the record conversion needs of as many applications and industries as possible, the method and device for converting or processing records between input and output devices must also be flexible. This is especially important where the number and types of input/output devices is likely to be significant.

This flexibility is preferably accomplished by utilizing multiple logical message interfaces (MI's), 28 and 30, which may easily be installed and supported in the transaction engine. For ease of explanation, only a single such input interface 28 and output interface 30 are depicted in Fig. 3B. Each logical message interface 28 and 30 will be unique to a record type to be input to or output from the conversion engine 23. Using different logical message interfaces for different record types allows a user to use only the interface(s) necessary to support a specific application.

All records input into the conversion engine 23 are converted into a format specific or "known" to the transaction engine i.e., the UDF. By converting all records into a specific format, a variety of different inputs and output devices can be installed, added, or deleted, without impacting the basic conversion functions of the transaction engine.

The data records received are temporarily stored in the super-record portion 23b of the conversion engine 23 for subsequent conversion to an output record.

All records to be output from the conversion engine 23 are first converted from the same "known" format i.e., the UDF to a desired output format. Again, by using a known format as the basis of conversion, different and multiple output record formats can be supported regardless of how the input record was formatted.

As further explained below, the conversion engine 23 processes records and/or messages bi-directionally so in addition to processing input records, application controls, e.g. acknowledgements, retransmission requests, flow control commands back to the generating devices 12 from the host device 20 may also be processed.

A library 23a supports conversion routines used to process either from an input record format to the universal format or from the universal format to the output format. By using a separate specialized processor for the conversion routines and by storing the conversion routines in the library, records will be processed most efficiently since only the conversions required will be performed, and the addition and deletion of

conversion routines is readily accomplished without impacting the entire transaction engine system.

As explained above, the interface that defines the type of record that will either be sent to or sent from the conversion engine 23 is the message interface (MI). The MI not only defines the type of input or output data being handled, it also establishes the conversion rules for a specific type of record, and the application controls that will be applied to the input or output records and/or messages.

The UDF is the definition of all fields supported within a specific conversion engine. For example, as illustrated in Fig. 3B a UDF may contain a date field that, by definition, always uses a Julian Date format. Received records that contain date fields may use a different format (e.g., MMDDYY), which are always converted into a Julian Date. Output records may also require a date field in a different format (e.g. DDMMYY), but the date would always be converted from a Julian Date format.

As alluded to above, the record that is temporarily stored in the transaction engine (before it is converted to an output record) is referred to as a "super-record" and is stored in the super-record 23b in Fig. 3B. This aspect of the conversion device will be further discussed below in connections with Figs. 9A, 9B, 11A and 11B. The super-record is comprised data fields in a format that is defined by the UDF.

Examples of conversion routines which may be found in the library 23a are an ASCII-to-EBCDIC routine, or a date conversion routine. For example, in Fig. 3B, a simple record consisting of date field and a time field is illustrated. Further, in the example of Fig. 3B, it is assumed that the input generating

device 12 requires an acknowledgement¹⁶ from the output host device 20 that a record has been received before it will send a second record. In such an example, the transaction engine operates as follows:

1) A data record from an input generating device 12 is received by the transaction engine 22. The record is defined as record type (A) based on the physical and logical port through which the record arrived as well as the content of the record itself. The rules for converting the fields within each record are issued by input MI 28.

2) The conversion routine library 23a is then called on to convert a date field and a time field. The routines in the library 23a will convert the date field and time field to the UDF format.

3) Using the definitions provided by the UDF, the date and time field are converted into and temporarily stored in the super-record 23b.

4) Subsequently, an output conversion process is initiated by the output MI 30. The output MI 30 defines the output record type in the illustrated example record type "B" and establishes the rules for converting to that record format.

5) The conversion routine library 23a is called on to make the date field and time field conversion.

6) Using the data record that has been temporarily stored in the super-record 23b, the conversion to the output record format is completed and sent to the host device 20.

7) Once the record is received and the output or host device 20 issues an acknowledgement, i.e., a return or application control, to the MI 30' (which may or may not be the

same as the MI 30) that a record¹⁷ was received and that an additional record may be sent.

8) The conversion engine 23 convert(s) the application control and passes it through to the MI 28'. In some instances, the application control record will not require conversion and may be passed directly to the MI 28'. The MI 28' may be the same or a different MI than the MI 28 used to convert the original input record.

9) Once the application control is received by the input device, the process repeats itself.

In contrast to prior processing systems, the transaction engine and data conversion process described above, operates in "real-time," occurring on the order of milliseconds. This enables the present invention to be implemented into virtually any preexisting data record collecting system and data processing system. In that regard, it is noted that system speed may vary and systems operating with process speeds in range of even several seconds or tens of seconds are within the scope of the invention. It is also within the scope of the present invention for certain records to be held or buffered and later dumped to a host device on command (usually an automatically generated timing command).

Table 1 depicts an example of a record format that may be supported by an MI. As explained above, the MI defines the record structure that it recognizes and the conversion rules it uses to transform a record. Typically, an MI will preferably support either input records or output records but not both. Input MI's, 28, (referred to below as "IMI's") transform a data record from the specified input format to the UDF. Output MI's,

30, (referred to below as "OMIs")¹⁸ typically transform a data record from the UDF to a specified output format(s). The MI functions are preferably reversible. For example, in connection with application control messages, e.g., an acknowledgment from a data processing unit to an input device, the MI 30 will operate as an IMI and the MI 28 as an OMI.

In Table 1, to the right of each message interface component is an example detailing typical contents of that message interface component.

Table 1 The Message Interface

The MI Components	
Unique Header	Date and Time Jan: 01, 1996::083013 Sequence Number 015 Physical/Logical Port . . 41::L0 Protocol SNA Encryption None Compression LZH
Transaction Record Type	TACACS
Transaction Record Sub-Type	1) IF Data Field 1 = 'A' AND Data Field 4 = 123.345.567.789 THEN use Transaction Record Structure 'A' 2) The value in header field 3 determines the sub-type *** (<n>) <definition rule>
Transaction Record Structure	Data Field 1 :: Integer 8 Bytes Data Field 2 :: ASCII Character 2 Bytes *** Data Field <n> :: <data type> <data size>
Conversion Rules	Data Field ⇒ Data Field 3 (Ø) Data Field ⇒ Data Field 1 (Julian Date ⇒ GMT) *** Data Field <x> ⇒ Data Field <y> (<conversion rule>))

The specific components of the message interface of Table 1 will now be further described.

Header

20

The conversion engine 23 of the transaction engine 22 preferably associates unique information, i.e., a header with each record it receives. This header will not only allow the conversion engine to track each record back to its origin, but also provide useful information for an audit trail. A typical header should preferably contain:

- 1) The date and time that the conversion device received the record;
- 2) The physical port of the conversion device that received the record;
- 3) The logical port associated with the record;
- 4) The communication protocol used to receive the record;
- 5) Whether the record is compressed or encrypted and, if so, the method used by the sender of the data.

Record Type

The MI record type preferably identifies the class of records that the MI processes, in this example TACACS. When the conversion engine recognizes an input record format, it routes the record to the appropriate MI for processing.

Record Sub-Type

The MI record sub-type is intended, where appropriate, to apportion the record type into sub-classes. If an MI specification requires record sub-types, the conversion engine 23 will preferably examine the record header to determine the sub-type of the input record.

Record Structure

21

The record structure defines the length of the record and specifies the field definitions, that is, field length and record type. The conversion engine 23 uses the record structure to properly map and convert the record fields. If the MI specification includes record sub-types, a separate record structure may exist for each sub-type.

Conversion Rules

The conversion rules instruct the conversion engine 23 how to convert a record to the desired universal format and vice versa. As exemplified in Table 1, the conversion rules may instruct the conversion engine 23 to perform tasks such as the following:

- (i) Copy field 3 in the input record to field 5 in the UDF record;
- (ii) Convert the input record time field from Julian Data to GMT;
- (iii) Convert the input record status field from ASCII to binary;
- (iv) Perform no conversion

In some cases, the conversion engine 23 will receive a compressed input record. In that case, the transaction engine will preferably examine the input record header to determine the compression algorithm and the length of the compressed record; and preferably, industry standard algorithms will be used to decompress an input record before processing and to compress an output record after processing.

22

In other cases, the transaction engine 22 may receive an encrypted input record. In that example, it will examine the record (and information in the header) to determine such things as the encryption algorithm used and the length of the encrypted record. Assuming the input record uses algorithms known to the transaction engine 22, the input record is decrypted before processing and then encrypted after processing to create the output record. In addition, the transaction engine 22 may participate in various other security architectures under the control of the general data processing system associated with an input device(s) and its companion host devices.

In some cases, input records may use a different character language than the output record(s). In such cases, the transaction engine 22 preferably converts between such characters, such as for example, packed decimal, ASCII, EBCDIC, and binary characters.

Some other useful functions of the transaction engine 22 will now be described.

Padding

Padding a field with blank characters guarantees that the length of the field will remain constant. The conversion engine 23 preferably right hand justifies and/or left hand justifies, as appropriate, using appropriate characters (e.g., zero's for numerical fields and blanks for character fields).

Truncation

23

Truncation is preferably used to ensure that record mapping does not inadvertently overwrite the contents of any other record fields when copying the contents of a record field.

"No Conversion" Option

The conversion engine 23 preferably recognizes when not to convert a particular field or an entire record. When combined with UDF definitions (which is a form of record mapping), the "no conversion" option creates three situations:

- 1) Map the record field and convert its contents;
- 2) Map the record field, but do not change its contents;
- 3) Do not map and do not change the contents of the record field.

The first two situations address the conversion or non-conversion of specific fields in the input record. In both situations, the transaction engine maps a field from the input record directly to the output record. The transaction engine preferably processes all records that fall into any of these situations. In the last situation, however, the transaction engine is not changing the input record, but is simply passing the input record to the output port. When an input record falls into the last situation, the conversion engine 23 does not convert the input record into the universal format.

The transaction engine provides the ability to redirect the records that IMI 28 receives to the records associated with the selected OMI(s) 30. As described in more detail below, each IMI 28 "listens" to a logical port associated with a physical port for input records in a logical protocol format specified during system configuration. When an IMI 28 receives an input record,

it converts the record to the universal format based upon the IMI conversion rules or using calls to the library 23a. Similarly, the OMI conversion rules are used to convert the universal format to the desired output format(s). In this sense, the conversion engine 23 may be envisioned as essentially operating as a "switch" that routes the converted records to the appropriate logical port(s) (and thereby to the physical port(s)) associated with the desired OMI(s).

In the following discussion, the flow of record and requests between the input generating device(s) 12, transaction engine 22 and output host 20 will be more fully explained in connection with a preferred embodiment depicted in Figs. 8, 9A, 9B, 10A, 10B, 11A and 11B. Turning first to Fig. 8, a simplified system having one input generating device and one output host device is depicted for the purposes of simplifying the explanation of how records and messages flow within the system. After the input generating device 12 establishes a connection to the transaction engine 22, the transaction engine is able to receive data records from the input generating device. As alluded to above, the transaction engine 22 passes each data record that it receives through a series of processes that convert it from a known but arbitrary input format to a desired output format. Upon completion of the conversion process, the transaction engine 22 establishes a connection to the output host device 20 and sends the data record, in the desired output format, to the host device.

In addition to processing data records sent from the input generating device 12 to the output host device 20, the transaction engine 22 also processes response messages, i.e.

25

control applications, from the output host device 20 to the transaction engine 22 and/or to the input generating device 12. In that regard, the transaction engine 22 preferably accommodates the following scenarios:

1. no response messages;
2. ignore response messages;
3. process response messages.

In the first scenario, the transaction engine 22 configuration does not provide for response messages transmitted from the output host device 20. After the converted transaction record is sent to the output host device 20, either the processing of that data record is complete or the processor can generate its own response to the input device. In the second and third scenarios, the configuration expects, for each data record, that the output host device 20 will send a response message to the transaction engine 22. In the second scenario, the transaction engine 22 will accept a response message from the output host device 20, but will not pass that message along to the input generating device 12. The third scenario accepts a response message from the output host device 20 and processes that message in a manner similar to an incoming conversion process applied to the input data record.

The data flow for a typical input data record from a record generating device will now be discussed in connection with Fig. 9A. The transaction engine 22 can be considered as comprising a series of transaction processors, TP1-TP6, some of which interface with the IMI's 28 (also see Fig. 3B), i.e., TP1 to TP4, for receiving and/or sending records/messages from/to an input generating device(s) 12 and some of which interface with

an OMI 30 (also see Fig. 3B) i.e., TP5 and TP6, for sending and/or receiving data records/messages to/from an output host device 20. Each of the transaction processors (TPs) will preferably interface with either an input record generating device 12 or an output host device 20 (see Fig. 3B). The transaction processors TP1 to TP4 that interface with the input record generating device(s) "listen" for a connection request. As soon as a connection is established to an associated transaction processor (TP) the transaction processor is available to process transaction records from the input generating device. When one of the transaction processors receives a transaction record, it copies that record to a shared memory 42. The transaction processors TP1 to TP4 access IMIs 28 and the transaction processors TP5 and TP6 access an OMI 30 each of which is executed using an associated dynamic link library ("DLL") 46. From the DLL 46 the transaction processor obtains information about a transaction record such as:

- transaction record type;
- transaction record length.

The transaction processors TP1-TP4 also access the DLL 46 of the IMIs 28 to convert that transaction record. The DLL 46 of the IMIs 28 applies a set of conversion functions - for example, converting time field in the data record from Greenwich Mean Time (GMT) to unix time t - to the transaction record data to translate it from one format to another. The transaction processors also manage the following data format conversions:

- input data format to the UDF (TP1-TP4);
- UDF to the output data format (TP5 and TP6).

When the data conversions are complete, an input side transaction processor executes a move signal which transfers the converted record to the appropriate output side transaction processor that issues a connection request to the output host device and subsequently transmits the converted transaction record thereto.

The shared memory 42 will vary the number of memory buffers 46 so as to "tune" overall system performance. Each available buffer 46 will contain several slots that track a data record conversion by providing a snapshot of the data at various points in the conversion process. When a transaction processor receives a data record, it copies the data in the record to the first buffer slot. After converting from the generating device format to UDF, the first buffer slot will contain the original data record and the second buffer slot will contain the UDF representation of that record. Similarly, after converting from the UDF to the output host device format, the second buffer slot will contain the UDF representation of the original data record and the third buffer slot will contain data in the output host device format. In addition, as alluded to above in the discussion of Fig. 3B, the shared memory 42 contains status information associated with the conversion process buffers 46 including the conversion mapping information, MI definitions, and the structure of the UDF.

The transaction processors examine the contents of the conversion mapping file and the MI definition file in a shared memory 42 to determine the protocol for processing a particular transaction record and/or message. In every case, the input record generating device needs to establish a connection to the

transaction engine 22. The transaction engine 22, in turn, must establish a connection with an output host device. An OMI definition file provides the address information for making the connection to the output host device. A conversion mapping file 50 also tells the appropriate transaction processor how to make the connection to the output host device. The conversion mapping file 50 tells the transaction processor how to process specific types of messages from the output host device. A data stream of messages may include several different types of messages. The conversion mapping file will tell the transaction processor whether it needs to convert a particular type of message, whether it can ignore a particular type of message, and, if it does need to convert that message, whether the message can be ignored, and if the message is to be stored, where it needs to store it once processing is complete. The MI definition 48 file in the shared memory 42 will also tell the transaction processor whether it should expect a response message from the output host device. If it does expect a response message from the output host device, the MI definition file 48 in the shared memory 42 will also tell the transaction processor whether the input record generating device needs to see that response message. If the input record generating device is not interested in seeing that message, the transaction engine 22 will simply discard any response messages from the output host.

Referring now to Fig. 9B, the flow of a response message, if any, from an output host device to an input record generating device will be described. As explained above, the transaction engine 22 may be configured to handle response messages from the

output host device. When such messages are received, the data flow will follow a path similar to but reversed from that described above in connection with the processing of transaction records as explained above in connection with Fig. 9A. In this scenario, after a transaction processor sends a converted transaction record to its associated output host device, it will wait for a response message from that host device. When and if it receives a response message, the transaction processor (in this case TP5 or TP6) copies that message into the memory buffers 44 of the shared memory 42 and "calls" the Dynamic Link Libraries ("DLLs") 52 of the associated OMI, 30 which in the response mode is functioning as an input MI, (in the described embodiment the MI's are preferably bidirectional) to convert that response message from the response format to the UDF and on the DLL 46 of the associated IMI 28 (which in the response mode is functioning as an output MI) to convert the response message from the UDF to the input generating device record format. After completion of the conversion process, the transaction processor associated with the host device sends a move signal to the transaction processor associated with the input record generating device. This move signal triggers that transaction processor to send the converted response message to the input record generating device.

Fig. 10A is similar to Fig. 9A but has been simplified for ease of understanding to illustrate the flow of a record from an input record generating device associated with a specific transaction processor, TP4, to a single output host associated with a specific transaction processor, TP6. The flow of data in Fig. 10A is essentially the same as that of Fig. 9A except that

there are no response messages or signals from the host device back to the input generating device routed through the transaction engine.

Similarly, Fig. 10B is similar to Fig. 9B but has been simplified to illustrate the flow of a single response message from an output host device to its associated transaction processor, TP6, and, via the transaction engine, to a specific input record generating device via its associated transaction processor TP4.

Referring now to Fig. 11A, a brief discussion of some of the processing protocols preferably employed in an exemplary embodiment of the transaction engine 22 will now be discussed. In this exemplary embodiment, the transaction engine 22 preferably uses multithreading to make the most efficient use of its central processing unit ("CPU"). As will be understood by one of many skill in the art, the transaction engine 22 may also employ multiprocessing techniques, i.e., plural CPUs. As is known in the art, multiprocessing takes advantage of a processing task that can be decomposed into several independent tasks. In multiprocessing each independent task is assigned to a separate CPU. This approach allows the independent tasks to be performed concurrently and speeds the overall execution time.

Similarly, an independent task can be further decomposed into two or more sub-tasks. As is also known to those skilled in the art, multithreading can be used to perform these sub-tasks concurrently to make the most efficient use of a single CPU. While one thread is reading an input port, another thread can perform a calculation, and a third thread can write to an output port.

3/

As further detailed below, the transaction engine 22 preferably uses multithreading in a way which associates a command queue with each thread. These command queues are the communication mechanism for the threads. Each transaction processor may comprise the following threads:

- a port listener (TP Listen);
- a session manager (TP Session Manager);
- a session translator (TP Translate);
- a session reader (TP Read);
- a session writer (TP Write).

Port Listener Thread

As explained above, the transaction engine 22 includes and controls a series of IMIs 28 and OMIs 30. Each MI has associated with it a logical port (32 and 34 respectively) and a DLL 46 (either discretely associated with a specific MI or as a distributed asset in the transaction engine 22) to convert the messages received by that port. The port listener thread (TP Listen 70) listens to the logical port 32 for an input record generating device requesting a new connection. When the listen thread receives a new connection request, it issues a session request to the TP Session Manager thread 72 and the transaction engine 22 creates a series of session threads to process the incoming data record. The TP Session Manager thread 72 accesses the associated transaction processor mapping file in the shared memory 42 to determine the IMI 28 that it should use. Similarly, the transaction processor accesses the IMI definition file in shared memory to determine the logical port 32 and address to listen to for the input generating device. When the

³²

TP Listen thread 70 receives a connection request, it issues a command to the TP Session Manager thread 72 command queue. This command triggers the TP Session Manager thread 72 to start up the TP Translate thread 74, TP Read thread 76, and TP Write thread 78.

Session Manager Thread

The TP Session Manager thread 72 manages the creation, administration, and deletion of all active transaction processor session threads (i.e., the session translator, session reader, session writer). The TP Session Manager thread 72 receives requests for new connections from the TP Listen thread 70. The transaction engine 22 allocates a series of session threads, sometimes referred to as an input session group, for each new connection. When the session threads are finished with the connection, the session manager creates a session remover thread which issues a command to each individual session thread and then clears the system. The TP Session Manager thread 72 may also issue the following commands to perform the following functions:

- retrieve the identity of the shared buffer element slot 44 that contains the outgoing record;
- determine whether an output session thread exist for the output record;
- create an output session thread(s) if it/they do not already exist;
- signal the output session threads when the record conversion process is complete;

- manage the allocation and deallocation of buffers 44 in shared memory 42.

When a session thread determines that a connection requires termination, either through natural completion or an error, the TP Session Manager thread 72 is notified through a remove session command with the session ID as an argument. The TP Session Manager thread 72 will then typically invoke a session remove thread and notify the related session threads to terminate by placing a command in their command queue. The session remove thread is a short-lived thread that methodically clears the threads which the TP Session Manager thread 72 originally created. The TP Write thread 78 is preferably the last session thread that the session remover thread eliminates. Prior to the TP Write thread 78 termination, it will send a disconnect over the connection to the downstream host device. Once the session threads are removed, the session remove thread clears the session information from the shared memory 42 and removes any extraneous record structure information.

Session Read Thread

The TP Read thread listens 76 to an input record generating device for a connection request. When the record arrives, the TP Read thread 76 queries the DLL 46 of the IMI 28 associated with the connection for the number of bytes that the records should contain. Once the TP Read thread 76 receives a complete message, it passes control to the TP Translate thread 74 by placing a command in the TP Translate thread 74 command queue.

Session Translate Thread

34

The TP Translate thread 74 receives messages from the TP Read thread 76. The associated DLL 46 of the MI 28 for the connection will be queried to determine the message type of the record. The message type is a search parameter in the conversion map file retained in shared memory 42. This file will inform the TP Translate thread 74 which of the MIs, 28 or 30, needs to receive a copy of this transaction record. For each interested output MI 30, the transaction record data is copied into a separate buffer slot 44 in the shared memory 42. Each of these buffer slots 44 goes through the conversion process from the input format to the UDF, and finally to the output format.

Session Write Thread

The TP Write thread 78 task will now be described in connection with Fig. 11B. The TP Write thread 78 writes the converted transaction record data to the port associated with the output host device. The TP Session Manager thread 72 signals the TP Write thread 78 when a transaction record conversion is complete. When the TP Write thread 78 receives this signal, it accesses the appropriate buffer slot 44 in shared memory 42 and writes that record to the associated output port, e.g., ports 34 or 34' (Figs. 4-7).

Switching

Turning now to the system embodiment of Fig. 4, an exemplary input record generating device 12 is depicted which is connected to an input port 32 (in this case a serial port) of the transaction engine 22. The conversion engine 23 passes

35

every input record that it receives over the input port 32 to the IMI 28 for conversion to the UDF. The conversion engine 23 senses when the conversion to the UDF is complete and signals the OMI 30 to convert the record to the desired output record format. The conversion engine 23 routes the converted record to the output port 34 (in this case a serial port). In this embodiment, the output port 32 is preferably connected to the output host device 20.

In the embodiment of Fig. 5, many-to-one or many-to-many switching based on addresses of the input devices is depicted. In this exemplary embodiment the input generating devices 12 are nodes of a network 20. Each input generating device 12 has a unique address. The transaction engine 22 receives input records from every networked input record generating device 12 through a single connection (e.g., an Ethernet port 32). When the transaction engine 22 receives an input record, based on its origin, it directs the appropriate IMI (28 or 28') to convert the record to the UDF. The conversion engine 23, using, e.g., the transaction processors described above, senses when the conversion to the universal format is complete and signals the appropriate OMI (30 or 30') to convert the record to the desired output format. The conversion engine 23 routes the converted record to the appropriate output port, 34 or 34'. In this embodiment, the output ports 34 and 34' are connected to the output host devices 20 and 20' respectively.

The embodiment of Fig. 6 is similar to the embodiment of Fig. 4 but depicts one-to-many switching based on information contained in the data or created by the input device. In this exemplary embodiment the input record generating device 12 is

36

directly connected to the port 32 of the transaction engine 22 (in this example, a serial port). The difference, however, is that the data record from this single input device 12 is distributed to multiple output host devices 20 and 20', in this example via the SNA protocol. The transaction engine 22 passes every data record that it receives from the port 32 to the IMI 28 for conversion, by means of the conversion engine 23, to the UDF. The conversion engine 23 (using, e.g., the transaction processors described above) senses when the conversion to the UDF is complete and signals the appropriate OMI (30 or 30') to convert the record to its particular output format. The conversion engine 23 then routes the converted records to the appropriate output port(s), 34 and/or 34', each of which is, in the illustrated embodiment, connected to an output host device, 20 and 20' respectively.

The embodiment of Fig. 7 depicts either many-to-many or many-to-one switching based on information in the data created by the input device. In this embodiment, a plurality, i.e., at least two, input record generating devices 12 and 12', are each connected, via respective input ports 32 and 32' to the IMI's 28 and 28' respectively. In this embodiment, the records from each of the input record generating devices 12 and 12' may be routed to one or more host devices 20 and 20'. The conversion engine 23 converts records received at either of the IMI's, 28 or 28'. The conversion engine 23, again using the transaction processors described above, senses when the conversion to the UDF is complete and then signals the appropriate OMI(s), 30 and/or 30' to convert the records to its particular output format. The conversion engine then routes the converted records to the

appropriate output port(s), 34 and/or 34' each of which, in the illustrated embodiment, is directly connected to an output host device, e.g., 20 and 20' respectively.

The foregoing description of the preferred embodiments of the invention have been presented for purposes of illustration and description. They are not intended to be exhaustive or to limit the invention to the precise forms disclosed, and obviously many modifications and variations are possible in light of the above teaching. These embodiments were chosen and described in order to best explain the principles of the invention and its practical application to thereby enable others skilled in the art to best utilize the invention in various embodiments and with various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the claims appended hereto.

We Claim:

1. A real time record retrieval and storage system comprising:

at least one record generating device for generating a record of a predetermined input format;

a conversion device including an input interface responsive to a signal from said record generating device for receiving the record from said record generating device, converting said record into a universal format, and from the universal format into a predetermined output format for an output host device; and

an output host device responsive to a signal from said conversion device for receiving data records in said predetermined output format.

2. The system according to claim 1 further comprising a plurality of record generating devices, wherein said conversion device further comprises an input interface responsive to each of the record generating devices.

3. The system according to claim 2 wherein said plurality of record generating devices are networked and said input interface comprises a network interface.

4. The system according to claim 1 wherein said conversion device comprises at least two input interfaces, each for supporting a different predetermined input record format.

5. The system according to claim 4 wherein each of said input data interfaces converts one of said predetermined input record formats into a universal format.

6. The system according to claim 5 wherein said conversion device further comprises at least one output

interface, said output interface being operable to convert records from said universal format into said predetermined output format.

7. The system according to claim 2 wherein said data conversion device further comprises at least one input port and at least one output port and wherein at least one of generating devices is directly connected to said input port; and wherein said output host device comprises a data processing device directly connected to said output port.

8. The system according to claim 3 wherein said conversion device further comprises an input port connected to at least two input interfaces, and wherein at least two of said plurality of record generating devices are configured as nodes on a network, each of said nodes having a unique address, and wherein said conversion device receives input records from said networked generating devices via said inlet port and, based on said unique address, directs said input records to one of said input interfaces.

9. The system according to claim 8 wherein said data conversion device further comprises at least two output interfaces, each of said outlet interfaces being connected to an output port and wherein said conversion device includes means for directing converted records to an appropriate one of said output interfaces, said system further comprising at least two output host devices for receiving output records, each of said output host devices being connected to one of said outlet ports.

10. The system according to claim 1 further comprising an input port wherein said input port is connected to said input interface, and wherein said record contains output host device

40

address data, said system further comprising a plurality of output host devices and said conversion device further comprising an output interface associated with each of at least some of said host devices and wherein the record from said record generating device is distributed to one or more of said some host devices in accordance with said address information.

11. The system according to claim 1 wherein said conversion device further comprises at least two output interfaces, each of said outlet interfaces being connected to an output port and wherein said conversion device includes means for directing converted records to an appropriate one of said output interfaces, said system further comprising at least two output host devices, each of said output host devices comprising a data processing devices hard wired to one of said outlet ports.

12. The system according to claim 1 further comprising a plurality of record generating devices for generating records of respective predetermined formats and wherein said conversion device comprises a port and a input interface for at least two of said plurality of generating devices whereby the records from said at least two record generating devices may be directed to one output host device.

13. The system according to claim 1 further comprising at a plurality of record generating devices for generating records of respective predetermined formats and wherein said conversion device comprises a port and a input interface for at least two said plurality of generating devices and said output host device comprises at least two output host devices and said conversion devices comprises a port and an output interface for each of

41

said at least two output host devices, whereby records from either of said at least two record generating devices may be selectively directed to either or both of said at least two output host devices.

14. The system according to claim 1 wherein said conversion device further comprises at least one input port and two output ports and at least two output interfaces, said input port being hard wired to said record generating device, each of said output interfaces being connected to one of said output ports, and wherein said conversion device includes means for directing records converted into a predetermined output format to an appropriate one of said output interfaces, said system further comprising at least two output host devices, each of said output host devices comprising a data processing device hard wired to one of said outlet ports.

15. The system according to claim 14 further comprising a plurality of networked record generating devices connected to said input port.

16. The system according to claim 1 wherein said conversion device comprises at least two input ports and at least two output ports, each of said input and output ports having an associated interface and wherein the interfaces associated with the input ports include means for converting said input records into a universal format and the interfaces associated with the output ports include means for converting the universal format into one or more predetermined output formats and directing said predetermined output format records, via said output interface, to one or more of said output ports.

42

17. The system according to claim 1 wherein said conversion device is located remotely from said host device in the vicinity of said record generating device.

18. The system according to claim 2 wherein said conversion device comprises one or more discrete devices located remotely from said host device, each of said discrete devices being associated with at least one of said plurality of generating devices.

19. The system according to claim 2 wherein said conversion device is centrally located in the vicinity of said output host device.

20. The system according to claim 1 wherein said input interface includes a memory for containing conversion functions for converting records from the predetermined input format to the universal format.

21. The system according to claim 4 wherein said input interfaces each includes a memory for containing conversion functions for converting a predetermined input format to the universal format.

22. The system according to claim 6 wherein said output interface includes a memory for containing function for converting records from the universal format to said predetermined output format.

23. The system according to claim 1 wherein said conversion device further comprises a memory for storing records during the conversion process.

24. The system according to claim 23 wherein said memory is a shared memory comprising a plurality of buffer memories, at least one of said buffer memories being used to store records in

the input format, one of said buffer memories being used to store records in the universal format and one of said buffer memories being used to store records in the output format.

25. The system according to claim 6 wherein said input interface and said output interface are bidirectional whereby records from said host device may be received at said output interface, converted by said conversion device into said universal format and from said universal format into the predetermined input format of said record generating device whereby records from said host device may be received by said record generating device.

26. A conversion device for receiving records in different formats from a plurality of record generating devices, said device comprising:

- (a) at least one input interface for receiving incoming records, said input interface supporting a predetermined input record format;
- (b) a conversion engine for receiving an input record from said input interface in said predetermined input record format and converting said input record into a predetermined universal format;
- (c) at least one output interface for receiving records in said universal format and converting from said universal format into a predetermined output format in order to support said predetermined output format.

27. The device according to claim 26 wherein said conversion device comprises a real-time conversion device.

28. The device according to claim 26 wherein said conversion device includes a first storage means for temporary

44

storage of records during conversion and a library of conversion rules for converting input records from one or more record generating devices which generate records in one or more predetermined input record formats into said universal format and for converting records from said universal format into one or more predetermined output record formats.

29. The device according to claim 26 wherein said input interface further comprises a memory for containing functions for converting from the incoming record format to the universal format.

30. The device according to claim 26 wherein said output interface further comprises a memory for containing functions for converting from the universal format to the predetermined output format.

31. The device according to claim 26 further comprising a library means for containing conversion protocols for converting records between an input data format and the universal format and between the universal format and a predetermined output format.

32. The device according to claim 26 wherein said conversion engine further comprises a memory including buffers for storing records in the predetermined input format, the universal format and the predetermined output format during record conversion.

33. The device according to claim 25 wherein said input interface and output interface are bidirectional interfaces, whereby a record may be received at said output interface and a converted record output from said input interface.

45

34. A real time method of handling records received in one or more input record formats into one or more output record formats using a conversion device comprising the steps of:

receiving records containing information organized in a predetermined input format from a record generating device;

converting said records into a universal format by means of said conversion device;

converting said records from said universal format into a predetermined output format by means of said conversion device;

transmitting the records in said predetermined output format to an output host device.

35. The method according to claim 34 further comprising the step of recognizing a connection request from said record generating device prior to the step of receiving.

36. The method according to claim 34 further comprising the step of issuing a connection request to the output host device prior to the step of transmitting.

37. The method according to claim 34 further comprising the step of receiving, at the conversion device, a response message from said output host device.

38. The method according to claim 34 further comprising the steps of converting said response message from the format in which it was received to said universal format and from said universal format to a format recognizable by said record generating device and transmitting said converted message to said record generating device.

46

39. The method according to claim 37 further comprising the steps of converting said response message from the format in which it was received to said universal format.

40. The method according to claim 34 wherein said conversion device includes a memory having a plurality of buffer memories, and wherein said step of converting further comprises:

copying the records received in said predetermined input format in a first buffer memory;

copying the converted records in said universal format in a second buffer memory; and

copying the converted records in said predetermined output format in a third buffer memory.

41. The method according to claim 34 further comprising the steps of adding an additional host device and said step transmitting further comprises transmitting said records in the output format of said additional host device.

42. The method according to claim 34 further comprising the steps of adding at least one additional output format and said step of transmitting further comprises transmitting said records in said additional output format.

43. The method according to claim 34 further comprising the steps of changing the output host device, said charged output host device accepting a different output format than said predetermined output format, and said step of transmitting further comprises transmitting said records in said different output format.

44. The method according to claim 34 wherein said records as received include no output host device address and wherein

said second step of addressing is based upon the content of said
record.

1/17

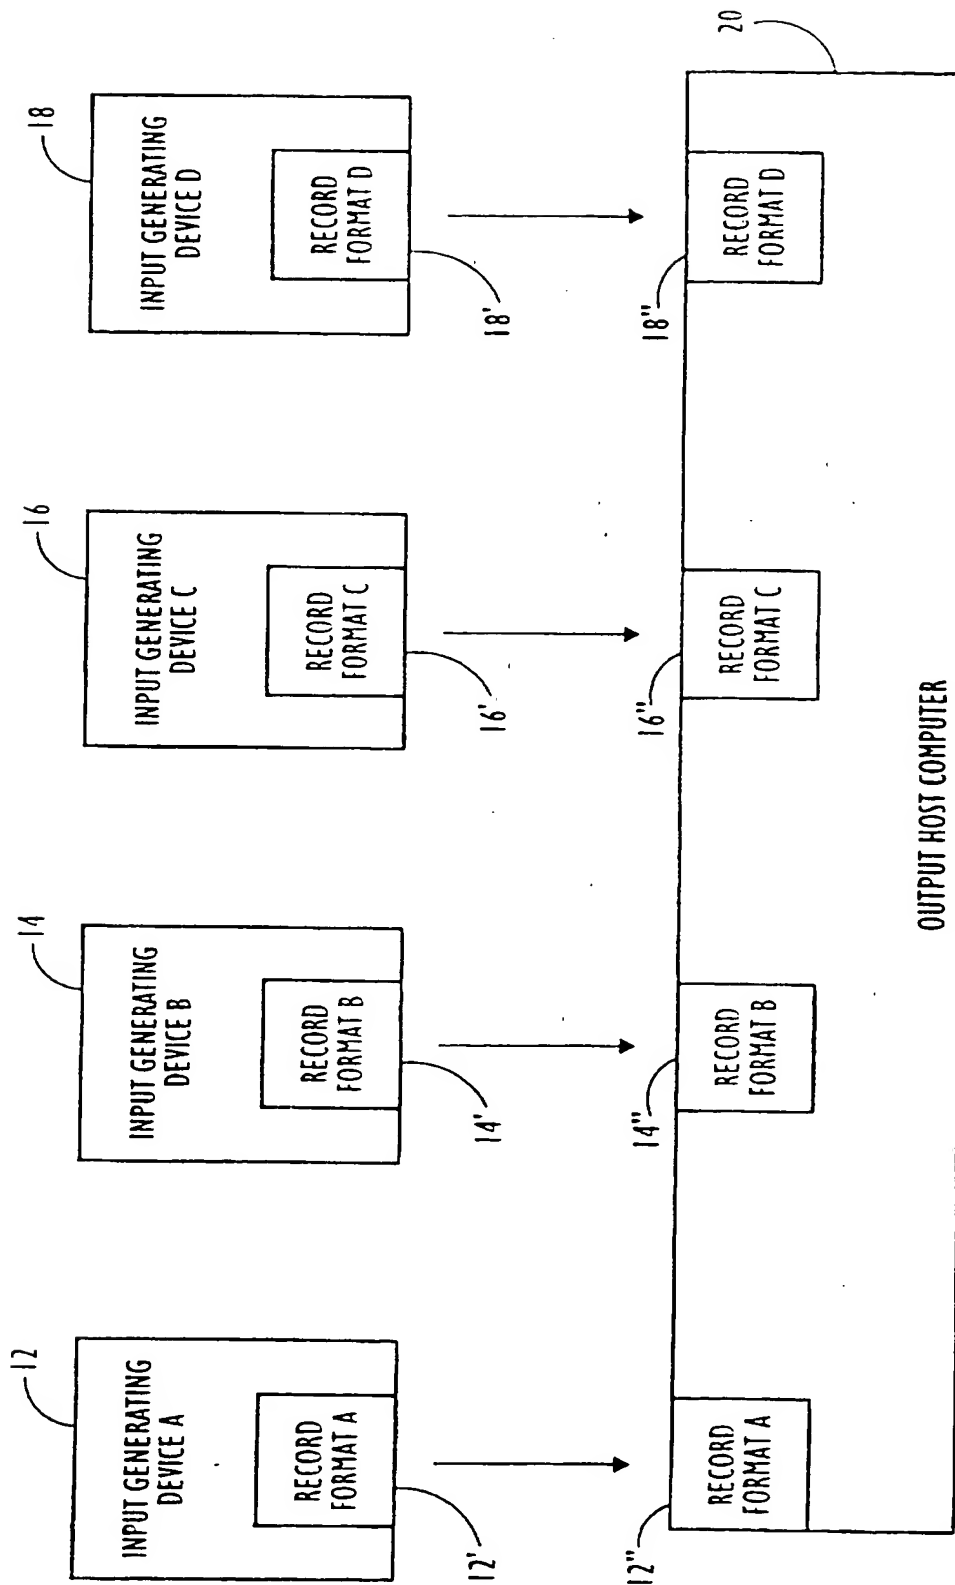


FIG. 1

2/17




FIG. 2A

3/17

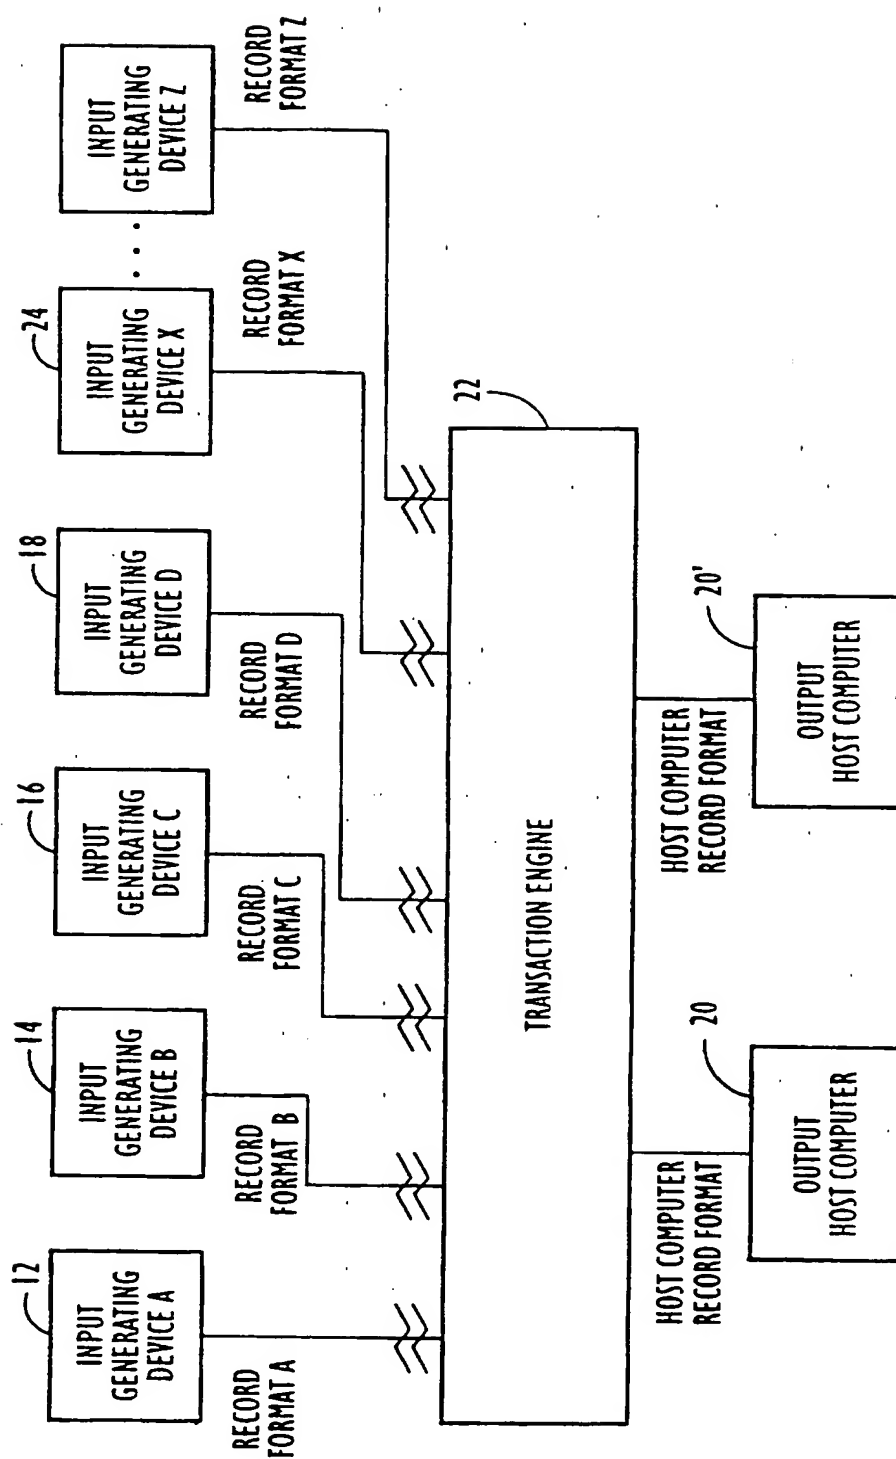


FIG. 2B

4/17

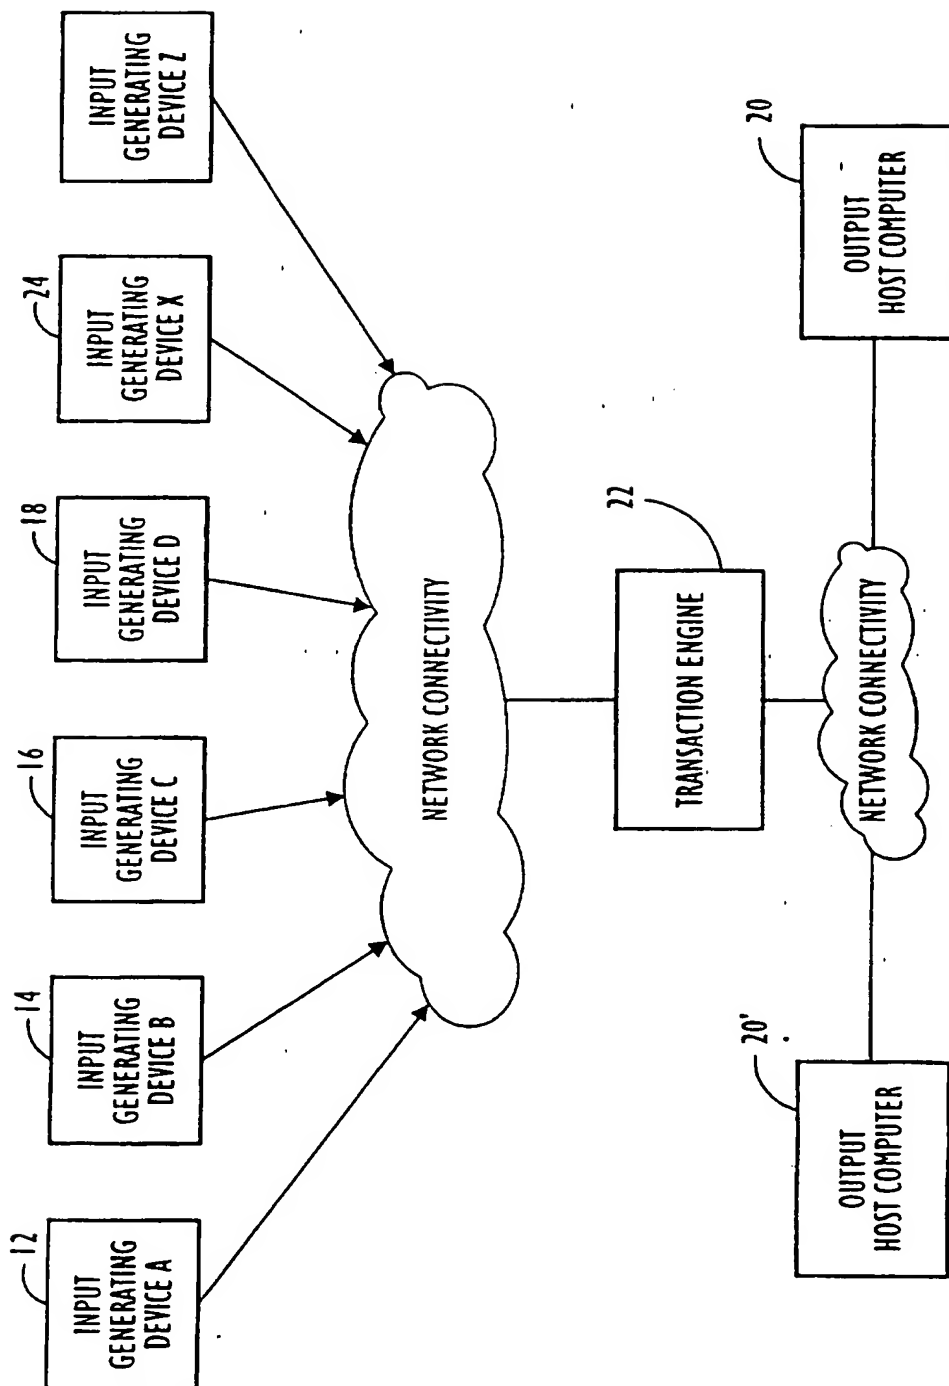


FIG. 2C

5/17




FIG. 3A

6/17




FIG. 3B

7/17




FIG. 4

8/n




FIG. 5

9/17




FIG. 6

10/17




FIG. 7




FIG. 8

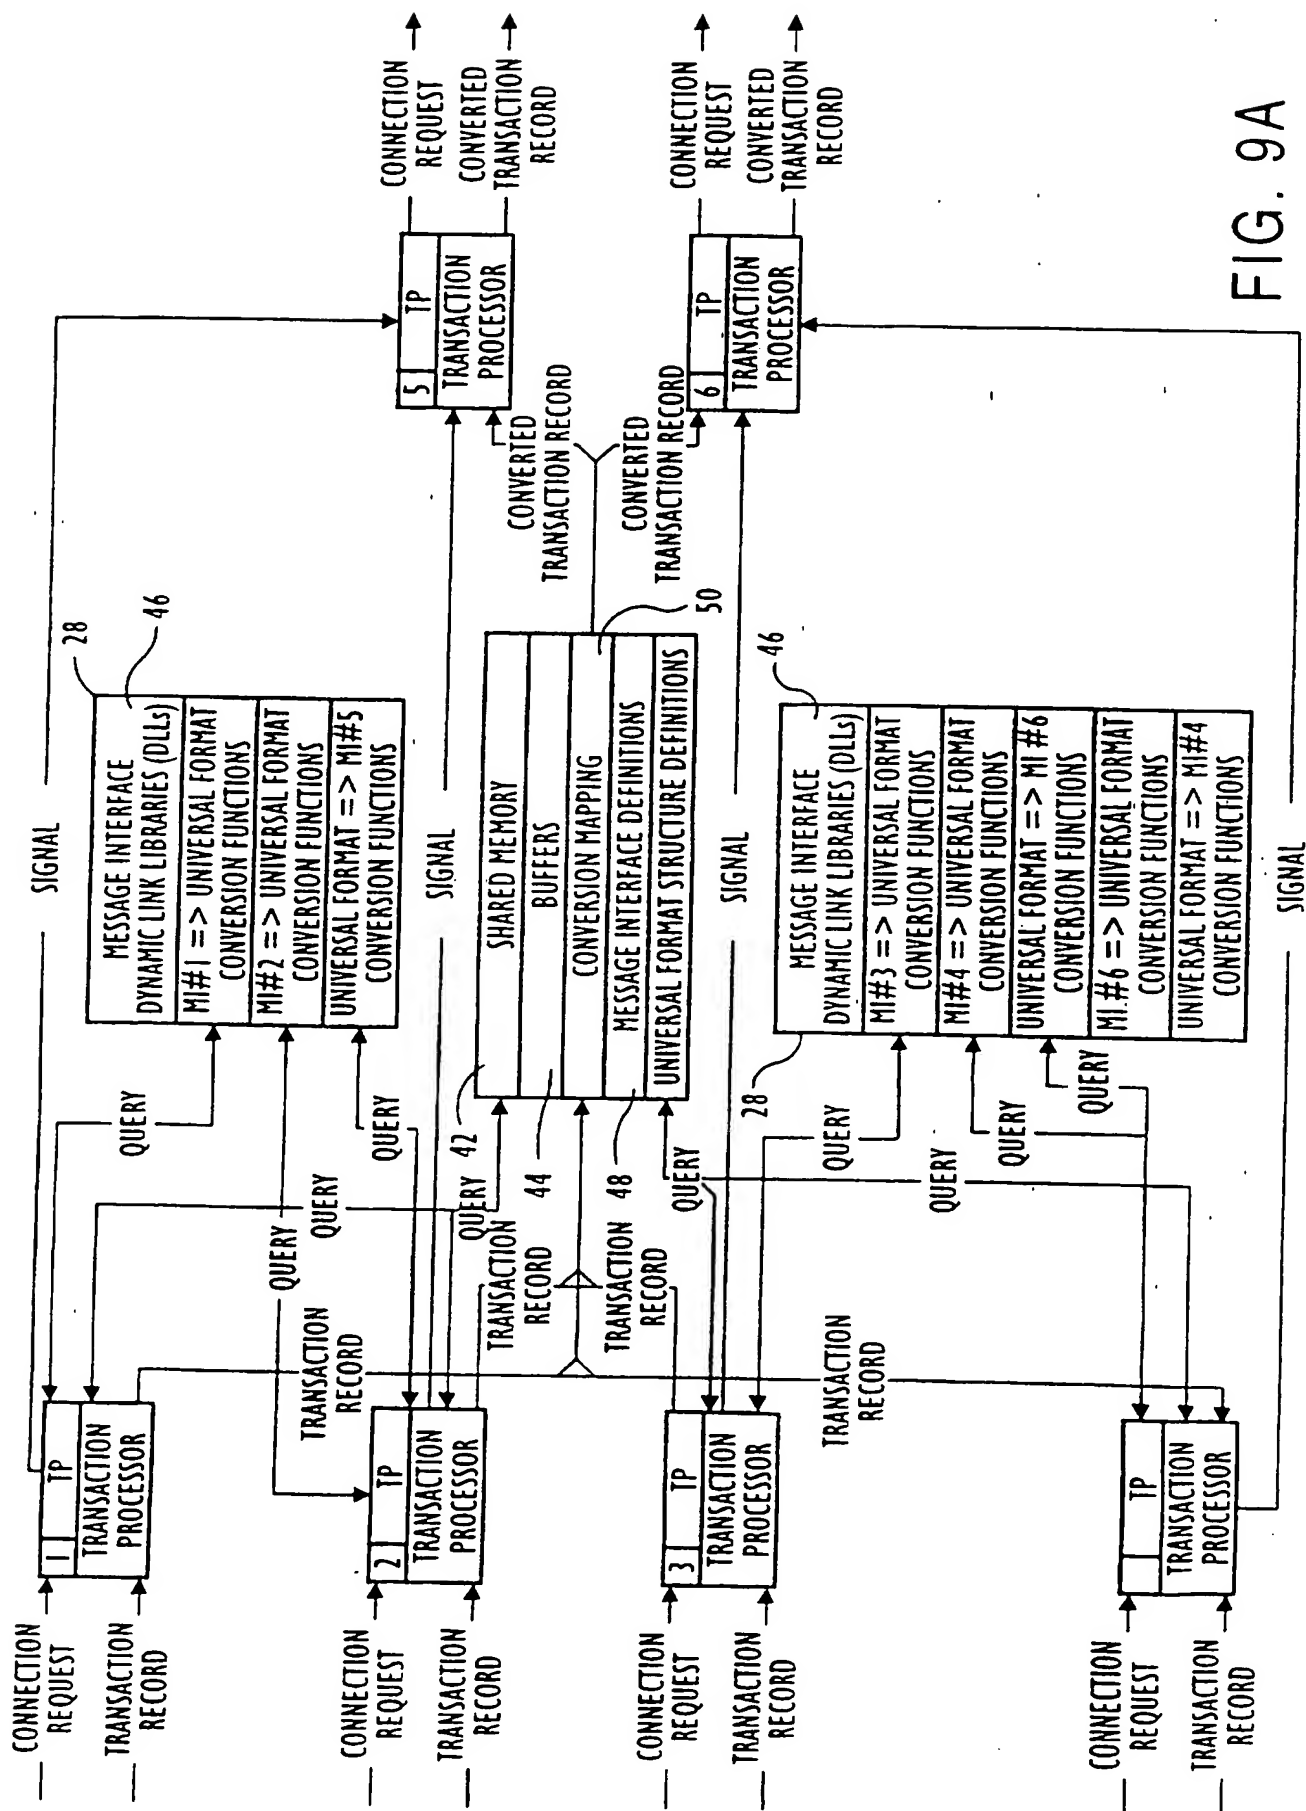


FIG. 9A

13/17

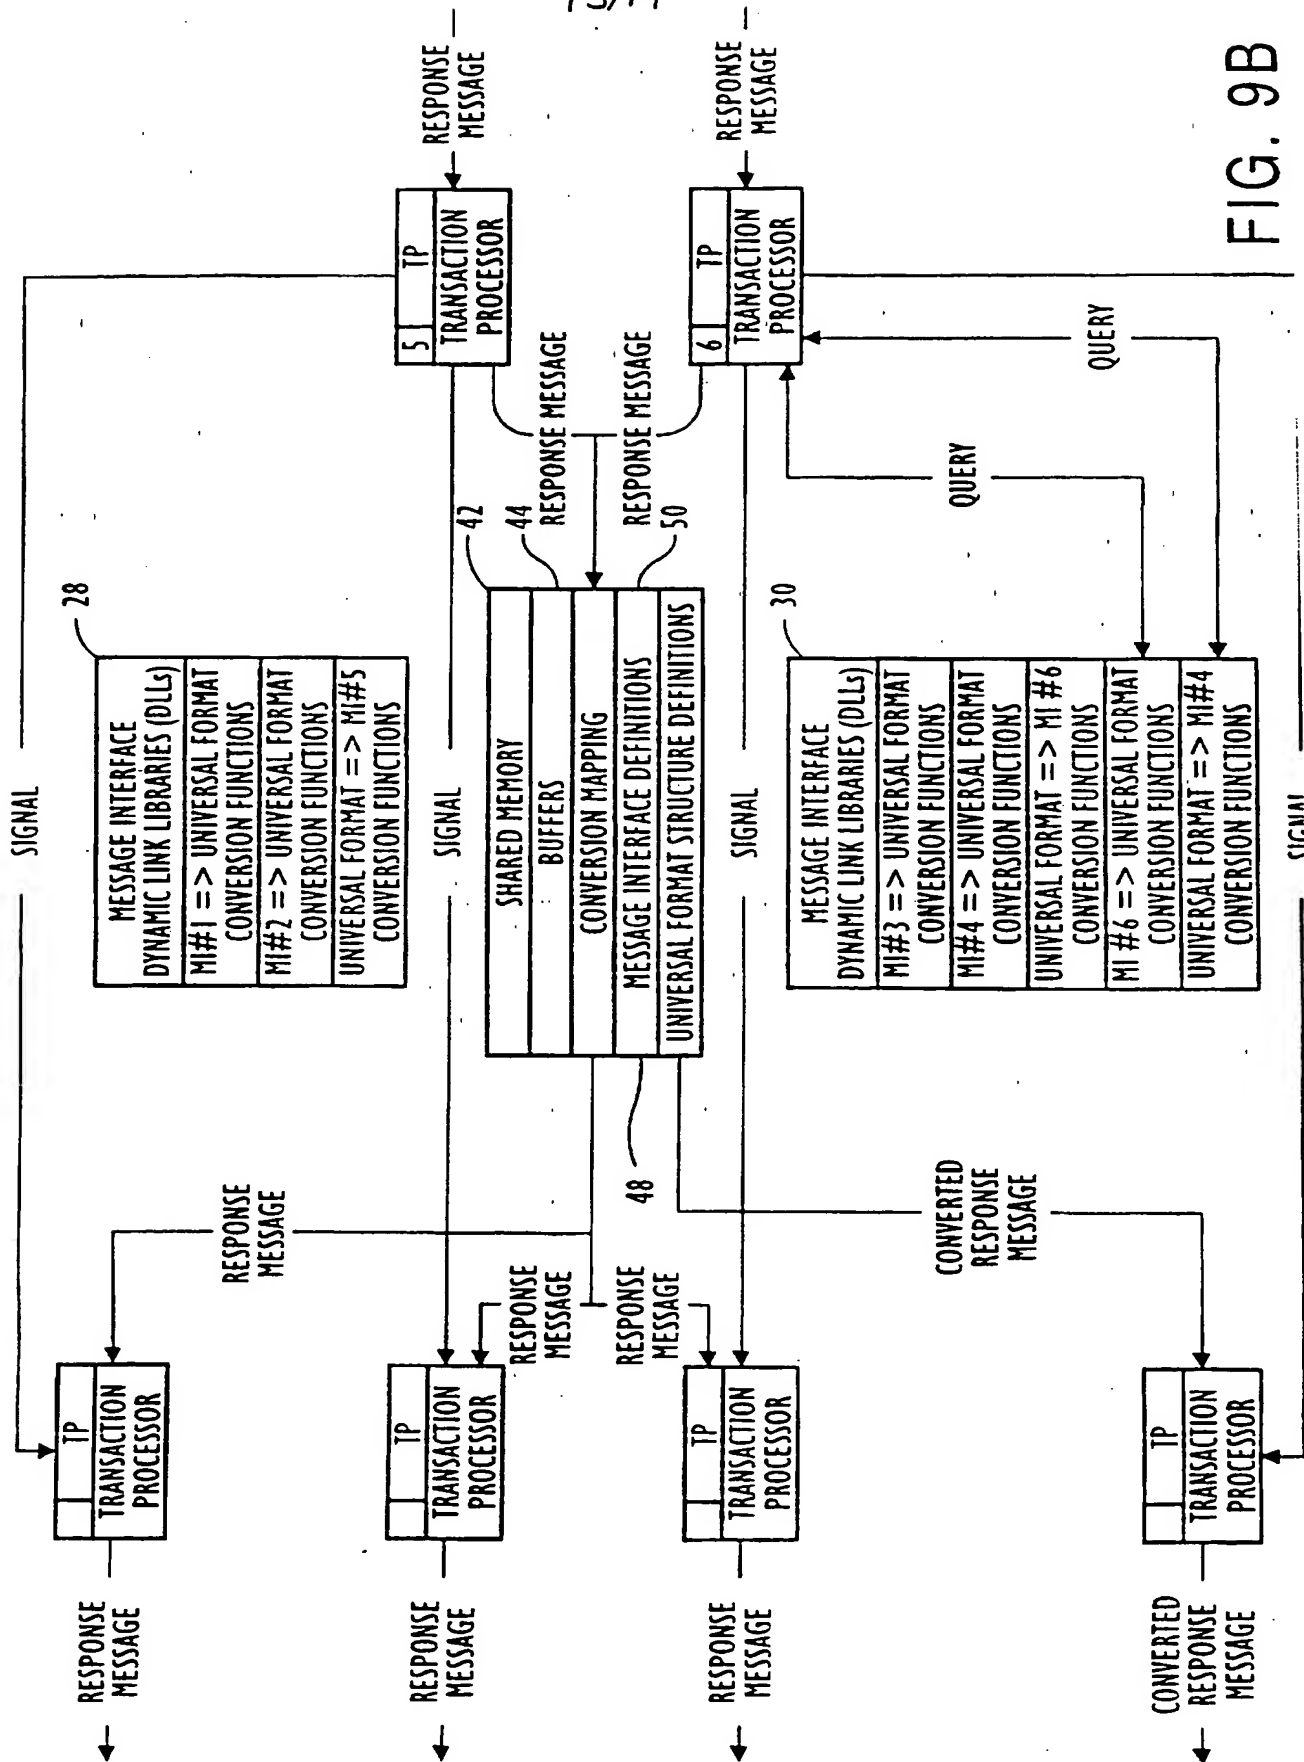


FIG. 9B

14/17




FIG. 10A

15/17




FIG. 10B





FIG. 11A





INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 17/30	A3	(11) International Publication Number: WO 97/40456
		(43) International Publication Date: 30 October 1997 (30.10.97)

(21) International Application Number: PCT/US97/06183

(22) International Filing Date: 24 April 1997 (24.04.97)

(30) Priority Data:

08/639,073	24 April 1996 (24.04.96)	US
08/825,276	27 March 1997 (27.03.97)	US

(71) Applicant: LINGUATEQ, INC. [US/US]; 1359 Beverly Road, McClean, VA 22101 (US).

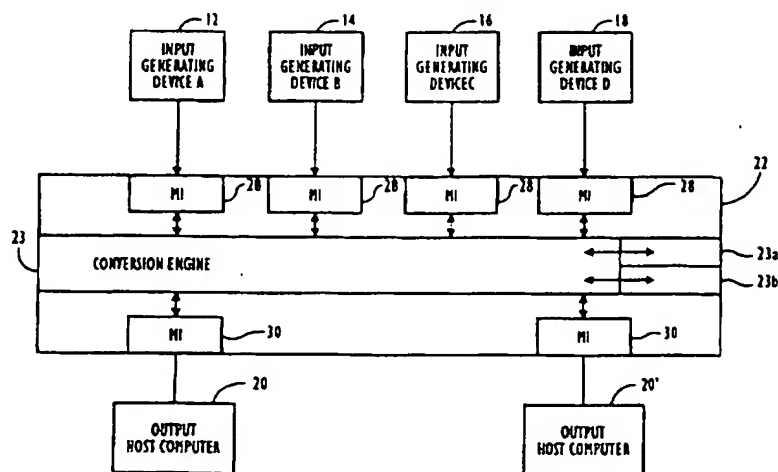
(72) Inventors: WILLIAMS, Dale, L.; 321 Springvale Road, Great Falls, VA 22066 (US). PETERSON, Michael, E.; 6821 Broyhill Street, McLean, VA 22101 (US).

(74) Agents: DEVINSKY, Paul et al.; McDermott, Will & Emery, 1850 K Street, N.W., Washington, DC 20006 (US).

(81) Designated States: AM, AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LT, LU, LV, MD, MG, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TT, UA, UG, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

Published*With international search report.**Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.*(88) Date of publication of the international search report:
4 December 1997 (04.12.97)

(54) Title: METHOD AND APPARATUS FOR RATIONALIZING DIFFERENT DATA FORMATS IN A DATA MANAGEMENT SYSTEM



(57) Abstract

A record conversion device and method, including a preferably software implemented conversion engine, which receives formatted input records from an input record generating device, converts the records into a different (universal) format in real time and, depending upon the output host device(s), into an output specific format and transmits the converted records to the output host device(s). Messages, i.e., information, from output host device(s) can also be transferred, via the conversion device, to the input record generating devices. By rapidly processing new and/or incompatible record formats, the invention minimizes the cost of maintaining information management systems and maximizes their usage. The conversion device input and output architecture is flexible and the conversion device can be used with multiple (networked or non-networked) input record generating devices and single or multiple record storage devices using multiple record output formats.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 97/06183

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X A	US 5 119 465 A (JACK MARTIN L ET AL) 2 June 1992 see column 1, line 50 - column 3, line 20	1,4-6,34 2,3, 7-33, 35-44
A	--- WO 93 23817 A (RELEASE MANAGEMENT SYSTEMS RMS ;TCS ENTERPRISES INC (US)) 25 November 1993 see page 6, line 14 - page 9, line 29; figures 7,8	1-44
A	--- "GENERIC DATABASE INTERFACE" RESEARCH DISCLOSURE, no. 318, 1 October 1990, page 794 XP000148635 see the whole document -----	1,26,34

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

B document member of the same patent family

Date of the actual completion of the international search

30 July 1997

Date of mailing of the international search report

22. 10. 97

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Fournier, C

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 97/06183

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5119465 A	02-06-92	NONE	
WO 9323817 A	25-11-93	AU 4237293 A US 5627972 A	13-12-93 06-05-97

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

This Page Blank (uspto)